

19.05.2025 Week 13 exercises: Distributed And Decentralized Machine Learning

Exercise 1 (Comparing Network Requirements in Federated Learning (FL) with Decentralized Learning (DL))

You run a FL server at home using the best Internet plan of your Internet Service Provider (ISP), which provides 1 Gbit/s download speed and 100 Mbit/s upload speed (see 2021 offerings by UPC as an example). Further suppose that:

- you are training a Deep Convolutional Network, such as LeNet [1], with 85098 parameters, with each parameter requiring a floating point value (32-bits) in memory; network latency is negligible;
- that the averaging time on the server is negligible;
- that the network, the server, and all clients are perfectly reliable and of similar performance;
- your ISP can fully and continuously guarantee the advertised download and upload speeds;
- your Federated Learning algorithm works in a synchronous fashion, i.e. all clients download the latest model parameters, then all clients update their parameters with a local gradient step, and then all clients upload the updated version of their model;
- performing a local gradient step takes 60 seconds, with no variance;
- that downloading or uploading the model parameters should take no more than 60 seconds (otherwise it would dominate over the cost of computation).

What is the maximum number of participants n_{max} you could support concurrently without exceeding the 60 seconds limit for either downloading or uploading parameters?

Exercise 2 (Economics of FL)

You are an Internet startup training the word prediction model of Google, in their Federated Learning paper (Section 3 “Large-scale LSTM experiments” in [3]) on a Cloud Provider. Compute the savings you could make in download costs if you had trained with DL instead.

For FL, assume the following:

1. parameters are represented as 32-bit floating point values
2. there are a total of 4,950,544 parameters
3. 200 clients participate in each round
4. training lasts 1000 rounds

5. you do not pay to send the models from participating clients to your Cloud Provider (ex: Amazon Cloud Pricing)
6. you pay $\frac{0.03\$}{2^{30}B}$ to download the updated model from the Cloud Provider to participating clients (ex: Amazon Cloud Pricing)
7. computation costs on the Cloud Provider are negligible.

For DL, assume you are not paying for bandwidth beyond your regular ISP subscription and you are using a local machine that you already own for computation.

Exercise 3 (Parameter server v/s All-reduce)

1. Describe how parameter server and all-reduce are two implementations of the same algorithm.
2. In the parameter server architecture, what is the bottleneck in terms of communication? What is the bottleneck in terms of computation?
3. In the all-reduce architecture, what is the bottleneck in terms of communication? What is the bottleneck in terms of computation?
4. Which architecture is more robust to failures? Why?

Exercise 4 (FL/DL aggregation)

Distributed learning systems rely on the aggregation of data produced by several computing nodes in order to iteratively produce improved versions of a machine learning model.

In the **federated learning** setting, a central server aggregates (averages) the local updates from the edge nodes (e.g., mobile phones) before sending the aggregated model back to a new set of nodes that will work on the next training iteration (i.e., all of them will train on the same new model aggregated by the server).

In the **decentralized learning** setting, on the other hand, there is no central aggregator. In order to compute aggregated models, nodes can only rely on updates shared by their neighbors. Thus, different nodes may train upon distinct models in a given training round. Before the next training iteration, this locally-trained model will be averaged with the updates received from neighbor.

Assume that the goal of a given distributed learning is to collectively achieve a global model that solves an image classification task using stochastic gradient descent (SGD) for optimizing the objective function F , i.e.,

$$\mathbf{x}_i^{(t;k)} \leftarrow \mathbf{x}_i^{(t;k-1)} - \eta \nabla F(\mathbf{x}_i^{(t;k-1)}; \mathbf{b})$$

where \mathbf{x}_i are the model parameters on node i , t is the epoch, k is the local step, η is the learning rate and \mathbf{b} is a mini-batch of training data.

Typically, the gradients ∇F are shared in the context of federated learning, whereas the parameters \mathbf{x}_i are shared in the context of decentralized learning. Why? Briefly describe how the choice between gradients or parameters is impacted by the distributed learning architecture, namely: (1) federated learning and (2) decentralized learning.

Exercise 5 (From Gossip to Learning)

You have seen how gossip can be used to compute the average of the attributes held by nodes in a network. Decentralized learning is doing something similar, but with a twist: instead of just averaging attributes, we are iteratively training and averaging attributes. **Can you modify the gossip algorithm that computes the average of the attributes to instead perform a variant of decentralized learning?** Clearly explain what the active and passive threads would do in your algorithm. Pinpoint the bottleneck(s) of your algorithm.

Exercise 6 (Decentralized Learning)

Decentralized parallel stochastic gradient descent (D-PSGD) [2] is an optimization algorithm often used for training models in DL. In each iteration of D-PSGD a node performs one or multiple steps of mini-batch SGD on their local model using their local data. Following this, the node transmits their local model to all direct neighbors in the graph. Upon receiving all models from all their neighbors a node aggregates the received models and their own into a new model that is used as their starting local model for the next iteration of D-PSGD.

There is a system containing 128 nodes organized in a graph with an average degree of 8. Assume that all nodes and connections between them are reliable and there are no stragglers. Furthermore, the distribution of each local dataset follows the distribution of the global dataset (IID data). The nodes perform 50 steps of D-PSGD, followed by one all-reduce round (assume naive all-reduce without any optimizations). What is the total number of messages exchanged in the given system during this training?

You have a global test set to evaluate each of the local models at different points during the described training process and measure their test loss. Take two arbitrary nodes from the graph and imagine how their respective test loss curves would look like during the training. What are some general conclusions you can come from from this imagination?

References

- [1] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [2] Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. *Advances in neural information processing systems*, 30, 2017.

- [3] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [4] Róbert Ormándi, István Hegedűs, and Márk Jelasity. Gossip learning with linear models on fully distributed data. *Concurrency and Computation: Practice and Experience*, 25(4):556–571, 2013.